

CONTINUATION APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: SPECTRALLY CONSTRAINED IMPULSE SHORTENING
FILTER FOR A DISCRETE MULTI-TONE RECEIVER

APPLICANT: GOPAL HARIKUMAR and DANIEL J. MARCHOK

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EF045063306US

June 27, 2003
Date of Deposit

SPECTRALLY CONSTRAINED IMPULSE SHORTENING FILTER FOR A
DISCRETE MULTI-TONE RECEIVER

5

Background

The invention relates to time-domain equalization in a discrete multi-tone (DMT) receiver.

Conventional single carrier modulation techniques translate data bits for transmission through a communication
10 channel by varying the amplitude and/or phase of a single sinusoidal carrier. By contrast, DMT, which is also referred to as Orthogonal Frequency Division Multiplexing (OFDM) or Multicarrier Modulation (MCM), employs a large number of sinusoidal subcarriers, e.g., 128 or 256
15 subcarriers. The available bandwidth of the communication channel is divided into subchannels and each subchannel communicates a part of the data. A DMT system may employ quadrature amplitude modulation (QAM) for each of the subcarriers.

20 OFDM-based systems transmit blocks of information bits. The time required to transmit one such block is called the symbol period. The time domain waveform that corresponds to one such block of bits is called a symbol.

Intersymbol interference (ISI) arises from the
25 characteristics of practical communication channels and limits the rate at which information can be transmitted through them. Specifically, communication channels typically have an Effective Discrete-Time Impulse Response (EDIR) that is greater than one sample time in length, which
30 causes ISI. ISI is a well-known phenomenon in single-carrier communication systems and there are many techniques for reducing it. The process of such ISI reduction is called equalization. ISI is discussed, for example, in Proakis, Digital Communications, McGraw Hill, 2nd Edition,
35 1989.

Equalization in OFDM-based systems is achieved by a two stage process. First, at the transmitter, an end-portion of each symbol is affixed to the beginning of the symbol to form what is called a Cyclic Prefix (CP). A
5 cyclic prefix that is greater than the EDIR of the channel prevents one symbol from interfering with another. Furthermore, it also facilitates a simple method of neutralizing the time-domain spread of each symbol forced by the channel. This is achieved by the use of a simple
10 frequency domain process in the receiver which requires one multiplication operation for each used subcarrier of the OFDM system.

The use of a Cyclic Prefix to reduce ISI is discussed, for example, in: Cimini, "Analysis and Simulation
15 of a Digital Mobile Channel using Orthogonal Frequency Division Multiplexing," IEEE Transactions on communications, pp 665-675 July 1985; Chow, "A Discrete Multi-Tone Transceiver System for HDSL applications," IEEE Journal on Selected Areas of Communications, 9(6):895-908, August 1991;
20 "DMT Group VDSL PMD Draft Standard Proposal," Technical Report, T1E1.4/96-329R2, ANSI 1997.

Another problem arising in conventional DMT systems is noise bleeding, which is the phenomenon of noise in one frequency band interfering with a signal whose subcarrier is
25 in another frequency band. Noise bleeding is caused, in general, by the Discrete Fourier Transform (DFT) operation at the receiver. Noise bleeding is discussed in, for example, Worthen et. al., "Simulation of VDSL Test Loops," Technical Report T1E1.4/97-288, ANSI 1997.

30 In a perfectly synchronized DMT system, the signal in one frequency band does not interfere with a signal whose subcarrier is in another frequency band. However, noise from one band may interfere with other less noisy bands and render them unusable. Techniques for dealing with noise-

bleeding include wavelet-based solutions. However, wavelet-based solutions are, in general, computationally intensive.

Summary

The invention provides a spectrally constrained
5 impulse shortening filter (SCISF) for use, for example, in DMT systems. The SCISF serves two primary functions.

First, the SCISF reduces intersymbol interference (ISI) by reducing the length of the effective discrete-time impulse response (EDIR) of the communication channel.
10 Conventional impulse shortening filters may have deep nulls in their frequency response. By contrast, the SCISF has a filter characteristic that is essentially free from undesired nulls that may attenuate or completely eliminate certain subcarriers.

15 Second, the SCISF reduces noise bleeding between subchannels by attenuating noisy channels in a manner that does not reduce the signal to noise ratio (SNR) in these channels, but reduces the noise power that may appear in the sidelobes of adjacent subchannels. The SCISF accomplishes
20 these functions by applying a frequency constraint to the signal based on a target spectral response.

In one general aspect, the invention features equalizing a channel in a multiple carrier communication system. The system includes a spectrally constrained
25 impulse shortening filter. Received noise power spectral density is measured. A target spectral response is computed based on the measured noise power. A frequency response for the spectrally constrained impulse shortening filter is selected based on the target spectral response. The
30 communication signal is filtered with the spectrally constrained impulse shortening filter.

Embodiments may include one or more of the following features. The noise power spectral density may be measured

at the output of the discrete Fourier transform. The spectrally constrained impulse shortening filter may be a time domain digital filter.

5 In another aspect, the invention features equalizing a channel in a multiple carrier communication system. The channel has an impulse response and is configured to receive a signal having a cyclic prefix. A target spectral response is computed. The impulse response of the channel is shortened so that a significant part of the energy of the
10 impulse response is confined to a region that is shorter than a target length, and the signal is filtered based on the target spectral response.

The target length may be a length of the cyclic prefix. The target spectral response may be computed from
15 measured noise power density, which may be measured at the output of a discrete Fourier transform. For example, the target spectral response may be the inverse of the measured noise power spectral density. The filtering step may be performed with a filter having a frequency response selected
20 to match the target spectral response. The shortening of the impulse response and/or the filtering may be performed by a time domain digital filter.

In another general aspect, the invention features selecting an impulse response for a spectrally constrained
25 impulse shortening filter in a multiple carrier communication system. Received noise power spectral density is measured. A cost function is computed using the noise power. The cost function is dependent on the impulse response. The dimensionality of a space over which the cost
30 function is defined is reduced and the cost function is minimized. The noise power spectral density may be measured at the output of a discrete Fourier transform. The cost function may be used to compute coefficients for the spectrally constrained impulse shortening filter.

In another general aspect, the invention features a spectrally constrained impulse shortening filter for a multiple carrier communication system. The system includes a channel that has an impulse response. A digital filter
5 structure receives the signal and apply a frequency characteristic to the signal. The frequency characteristic is determined by filter coefficients. Filter coefficients are selected to shorten the impulse response of the channel so that a significant part of the energy of the impulse
10 response is confined to a region that is shorter than a target length and to apply a frequency characteristic to the signal based on a target spectral response. The selected filter coefficients are input to the taps of the filter.

In another general aspect, the invention features a
15 receiver for receiving a multiple carrier signal from a communication channel having an impulse response. An analog-to-digital converter receives the signal from the communication channel. A spectrally constrained impulse shortening filter receives the signal from the analog-to-
20 digital converter and shortens the impulse response of the channel so that a significant part of an energy of the impulse response is confined to a region that is shorter than a target length. The filter also applies a frequency characteristic to the signal based on a target spectral
25 response. A discrete Fourier transform receives the output of the spectrally constrained impulse shortening filter and a decoder receives outputs of the discrete Fourier transform.

In another aspect, the invention features a modem in
30 which an encoder receives digital data and outputs a constellation of complex values. An inverse discrete Fourier transform receives the constellation from the encoder. A digital-to-analog converter is connected to the inverse discrete Fourier transform and outputs a signal to a

communication channel. An analog-to-digital converter receives the signal from the communication channel. A spectrally constrained impulse shortening filter shortens an impulse response of the channel so that a significant part of an energy of the impulse response is confined to a region that is shorter than a target length. The signal is filtered based on a target spectral response. A discrete Fourier transform connected to the filter and a decoder is connected to the discrete Fourier transform and outputs digital data.

In another aspect, the invention features software for causing a processor in a communication system to perform the following operations: measure received noise power spectral density; and compute a target spectral response based on the measured noise power. The software may also include instructions for causing a processor to compute filter coefficients based on the target spectral response.

In another aspect, the invention features software for causing a processor in a communication system to measure received noise power spectral density. A cost function is computed using the noise power. The cost function is dependent on an impulse response of a spectrally constrained impulse shortening filter. The dimensionality of a space over which the cost function is defined is reduced, and the cost function is minimized.

The techniques described here are not limited to any particular hardware or software configuration. They may find applicability in any computing or processing environment that may be used for a communication system. The techniques may be implemented in hardware or software, or a combination of the two. Preferably, the techniques are implemented in computer programs executing on a digital signal processor that includes a processor and a storage

medium readable by the processor (including volatile and non-volatile memory).

Other features and advantages will be apparent from the following description, including the drawings and the
5 claims.

Brief Description of the Drawings

Fig. 1 is a block diagram of a discrete multi-tone communication system.

Fig. 2 is a plot of effective discrete-time impulse
10 response (EDIR) of a communication channel including transmit and receive filters.

Fig. 3 is a plot of the shortened EDIR due to a spectrally constrained impulse shortening filter (SCISF).

Fig. 4 is a block diagram of a SCISF.

15 Fig. 5 is a plot of transmit signal power, signal power at SCISF input and noise power at SCISF input.

Fig. 6 is a plot of signal and noise power at the output of the SCISF.

20 Fig. 7 shows the filter response of a discrete Fourier transform for one frequency band.

Fig. 8 is a plot of the target spectral response of the SCISF, $G_d(\omega)$, versus the actual frequency response, $G(\omega)$.

25 Fig. 9 is a plot of signal-to-noise ratio at the output of the SCISF, the output of the DFT with the SCISF and the output of the DFT without the SCISF.

Description

As shown in Fig. 1, a discrete multi-tone (DMT) communication system 10 has a transmitter 12 and a receiver
30 14. The transmitter 12 accepts an input data bit stream which passes through a constellation encoder 20. The encoder 20 divides the serial input bit stream into blocks

of data. These blocks of data are further subdivided into smaller blocks corresponding to subchannels. Each of these smaller blocks are used to compute a complex value representing a constellation point. Each constellation
5 point corresponds to a subsymbol. The subsymbols are then output by the encoder. Taken together, the subsymbols constitute a symbol.

The subsymbols are supplied to an inverse discrete Fourier transform (IDFT) 30, which may be implemented, for
10 example, in a digital signal processor. The IDFT 30 outputs N time samples of a symbol. The time samples are processed by a parallel to serial converter 40 to form a single stream of time samples.

Following the parallel to serial converter 40, a
15 prefixer 50 adds a cyclic prefix to the beginning of each symbol to reduce intersymbol interference (ISI). Alternatively, the cyclic prefix may be added in the parallel to serial converter. After the cyclic prefix is added, the resulting signal passes through a digital-to-
20 analog (D/A) converter 60 for transmission to the receiver 14 through a communication channel 70. An analog transmit filter 65 may be included following the D/A converter to band limit the transmitted signal.

At the receiver 14, the signal passes through an
25 analog-to-digital (A/D) converter 75 and then through an impulse shortening filter 90. An analog receive filter 85 may be included prior to the A/D converter in order to band limit the received signal. A prefix stripper 100 strips the cyclic prefixes from the resulting symbols and a serial to
30 parallel converter 110 divides the stream of time samples into parallel signal paths that form the inputs to a discrete Fourier transform (DFT) 120. The DFT 120 converts the time samples into subsymbols. A decoder 130 converts

the subsymbols into a data bits and outputs the resulting data.

As discussed above, a cyclic prefix is added to each symbol prior to transmission through the communication channel to reduce the effects of ISI. The cyclic prefix is formed by copying the last ν time samples from the end of a symbol and placing them at the beginning of the symbol. To eliminate ISI, the length of the cyclic prefix, ν , is chosen to be longer than the effective discrete-time impulse response (EDIR) of the channel. However, because the cyclic prefix constitutes redundant data, increasing the length of the cyclic prefix reduces the efficiency of the communication system. For example, in a system having N time samples per symbol and a cyclic prefix of ν time samples, the efficiency of the system will be reduced by $N/(N+\nu)$. To maximize efficiency it is necessary either to minimize ν or to maximize N . However, increasing N increases the complexity, latency and computational requirements of the system and at some point becomes impractical. Accordingly, it is desirable to minimize ν .

An impulse shortening filter having an impulse response, $g(n)$, may be employed in the receiver to minimize the length of the cyclic prefix by decreasing the EDIR, $g(n)*h(n)$, of the effective communication channel, which includes the transmit and receive filters, the impulse shortening filter and the physical transmission channel. The use of an impulse shortening filter is referred to as time domain equalization. Decreasing the EDIR allows a shorter cyclic prefix to be used without increasing ISI.

Fig. 2 is a plot of the EDIR for a DMT test configuration having a communication channel that is 4500 feet in length and operates at a sampling frequency of 11.04 MHz (test loop 4, as described in "Very-high Speed Digital Subscriber Lines: System Requirements," Technical

Report T1E1.4/97-131R1, ANSI 1998). The EDIR includes the effects of a transmit filter, the communication channel and a receive filter. Fig. 3 shows the impulse response as shortened by the addition of an impulse shortening filter.

5 The impulse shortening filter is selected so that a significant part of the energy of the joint impulse response of the filter and the effective communication channel, $g(n)*h(n)$, is confined to a region that is shorter in length than the length of the cyclic prefix. Some prior algorithms
10 for computing $g(n)$ considered only shortening the EDIR and did not consider the spectral characteristics of the resulting impulse shortening filter. Such filters often had deep nulls in some frequency bands, which rendered some of the corresponding subchannels useless.

15 Since increasing the length of the cyclic prefix reduces system efficiency, the receiver may dynamically compute an optimal length for the cyclic prefix and may send that information to the transmitter. For example, the receiver may compute a set of impulse responses for the
20 impulse shortening filter based on a set of predetermined cyclic prefix lengths. The receiver then computes the system throughput for each particular cyclic prefix length. The length that maximizes system throughput is selected and the result of that selection is communicated to the
25 transmitter. The transmitter then operates using the selected cyclic prefix length.

 To avoid the possible attenuation of frequency bands, the spectral response of the impulse shortening filter is further required to have a spectral response that
30 $|G(\omega)|$, meets a specified spectral constraint. A spectral constraint of the form $|G(\omega)H(\omega)| > \tau$, where τ is a threshold, is sufficient to avoid nulls in the frequency response of the impulse shortening filter. However, it is possible to compute a spectral constraint or target spectral response,

$|G_d(\omega)|$, that provides additional performance improvements, such as reducing noise bleeding between subchannels. A spectrally constrained impulse filter is configured to have a spectral response that approximates the target spectral
5 response.

As shown in Fig. 4, the spectrally constrained impulse shortening filter (SCISF) 90 may be implemented as a time domain digital filter, which has a digital filter structure 210 with a number of taps 220 or filter
10 coefficient inputs for adjusting the filter response. The coefficients may be computed and supplied to the taps by a digital signal processor (DSP). Alternatively, the SCISF may be implemented entirely in software, i.e., within a DSP.

A desired spectral response may be applied to a
15 received signal using a filter that is separate from the impulse shortening or time domain equalization (TEQ) filter. For example, an analog filter may be placed prior to the A/D converter. However, the adjustability of such a filter would be limited. As a further example, a digital filter
20 could be added prior to the TEQ filter. Both of these configurations suffer the disadvantage that the TEQ filter may distort the desired spectral characteristics of the added filter. A filter also might be positioned after the TEQ filter, which would reduce noise bleeding, but might
25 reduce the impulse shortening provided by the TEQ filter. Accordingly, the SCISF integrates the TEQ (i.e., impulse shortening) function with the desired spectral response in a single filter.

In summary, the filter characteristic, $g(n)$, of the
30 SCISF satisfies two conditions. First, the effective length of the convolution of the filter characteristic with the impulse response of the communication channel, $g(n)*h(n)$, is less than a target length. Second, the cost function (error

function) between the target spectral response, $G_d(\omega)$, and the actual filter spectral response, $G(\omega)$, is minimized.

The target spectral response is an ideal filter characteristic that is selected to maximize the data bit throughput in the subchannels of a DMT system by reducing the impact of noise. There are many sources for noise in a DMT communication system, such as near-end cross-talk (NEXT), radio frequency interference (RFI) and noise generated in the communication channel (white noise). As shown in Figs. 5 and 6, the noise spectral density is generally not uniform across the frequency band of the communication system. This non-uniformity contributes to the problem of noise bleeding, in which noise in one frequency band interfering with a signal in another frequency band.

In general, noise bleeding is caused by sidelobes of filters in a DFT. Fig. 7 shows the filter response of a DFT for one frequency band or bin (i.e., bin 128). The first sidelobes (250) are only 13 dB below the main lobe (260). Therefore, noise located outside of bin 128, but within the first sidelobe of bin 128, i.e., approximately mid-way between bins 126 and 127, would appear in bin 128 with an attenuation of only 13 dB. Consequently, noisy subchannels in a DMT system may degrade the performance of non-noisy subchannels.

The target spectral response is essentially a spectral constraint that attenuates noisy channels more than non-noisy channels. The signal and noise in the noisy channels are reduced equally, so the attenuation does not affect the signal-to-noise ratio in these channels. However, because the absolute noise level in the noisy channels is reduced, there is less noise available to enter the sidelobes of the non-noisy channels. Hence, the noise bleeding problem is minimized.

To determine the target spectral response, the noise power spectral density (noise PSD) at the receiver must be known. The noise PSD may be determined, for example, by performing a periodogram on received data. This measurement is more complicated if a transmitter is transmitting, since the transmitted signal must be separated from the noise measurement. The noise PSD is determined by: (i) slicing the received constellation of subcarriers after the DFT to determine the nearest valid constellation point; (ii) determining an error signal based on the difference between the received constellation point and the valid constellation point; (iii) performing an IDFT on the error signal; and (iv) generating a periodogram (with windowing) from the error signals. The noise PSD may then be determined from the periodogram.

An example of a noise PSD characteristic for a DMT communication system is shown in Fig. 5 (test loop 4, as described in "Very-high Speed Digital Subscriber Lines: System Requirements," Technical Report T1E1.4/97-131R1, ANSI 1998). The transmit signal power is measured at the output of the transmit filter in the transmitter. The signal and noise PSD plots shown in Fig. 5 are measured at the input of the A/D converter in the receiver, which is prior to the SCISF.

Measured noise PSD is used by a digital signal processor (DSP) to compute a target spectral response, $G_d(\omega)$, using the algorithm described below. Alternatively, the inverse of the noise PSD may be used as an approximation for the target spectral response. A spectral response, $G(\omega)$, is then determined for the SCISF that minimizes the error between the spectral response of the SCISF and the target spectral response. A set of filter coefficients may then be generated to configure the SCISF to the determined characteristic. These calculations may be done periodically

to adjust the performance of the communication system. Frequency domain equalization coefficients and symbol synchronization may also be adjusted based on these calculations.

5 Fig. 8 is a plot of the target spectral response of the SCISF, $G_d(\omega)$, versus the actual frequency response, $G(\omega)$. The difference between the responses is only a few dB. Fig. 6 shows the signal and noise PSD at the output of the SCISF.

10 Fig. 9 shows the dramatic effect of the SCISF on system performance. Without the SCISF (i.e., using a filter that provides only impulse shortening), the signal-to-noise ratio (SNR) decreases significantly at the output of the Fourier transform (i.e., FFT or DFT) to less than about
15 7 dB. This decrease is due, in large part, to noise bleeding caused by the sidelobes of the Fourier transform. By contrast, with the SCISF, the SNR at the output of the Fourier transform tracks the SNR at the output of the SCISF within a few dB. Overall, the SCISF provides an improvement
20 in SNR.

 The target spectral response of the SCISF, $G_d(\omega)$ and the actual frequency response, $G(\omega)$, are derived from an energy constraint for the SCISF, i.e., the target spectral response must localize the energy of the effective impulse
25 response within a desired frequency band. The energy constraint is combined with a desired spectral response based on the measured noise power spectral density. The resulting cost function (or error function) is then minimized to obtain a practical filter characteristic for
30 the SCISF. This process is presented in detail below.

In the following derivation, all vectors are column vectors by default. Vectors are denoted by bold lower case letters (e.g., \mathbf{t}). The size m_t of a vector \mathbf{t} , is written $t(m_t)$. The components of a vector are denoted by lower case letters, e.g.,

5 $\mathbf{t}(m_t) = [t_0 \cdots t_{m_t-1}]^T$. The convolution of length $m_t + m_h - 1$ of vectors $\mathbf{t}(m_t)$ and $\mathbf{h}(m_h)$ is denoted $\mathbf{t} * \mathbf{h}$. The Toeplitz matrix of a vector $\mathbf{t}(m_t)$, $\mathbf{t}_{(m:n)} \triangleq [t_m \cdots t_n]^T$, is written $\mathbf{T}_{(m:n)}^{m_t}$, and $\mathbf{T}_{(m:n)}^{m_t} \mathbf{x} = (\mathbf{t} * \mathbf{x})_{(m:n)}$. The Discrete Time Fourier Transform (DTFT) of a vector \mathbf{t} may be represented as:

$$10 \quad T(\omega) = \sum_{n=0}^{m_t-1} t_n e^{-j\omega n}. \quad (1)$$

The EDIR of the channel and the impulse response of the SCISF are expressed as $\mathbf{h}(m_h)$ and $\mathbf{g}(m_g)$, respectively. The cyclic prefix has a length of m_c per symbol. The energy of the effective impulse response must be localized to a contiguous
15 region of target length m_1 , where $m_1 < m_c$, while satisfying a spectral constraint. The energy criterion to be satisfied by the SCISF may be written as:

$$\frac{\|(\mathbf{g} * \mathbf{h})_{(m:m+m_1-1)}\|^2}{\|(\mathbf{g} * \mathbf{h})\|^2} \geq \alpha \quad (2)$$

for some $0 \leq m \leq m_g + m_h - 2$ and some $0 < \alpha < 1$. Defining the set S^m as

$$20 \quad S^m = \left\{ \mathbf{g} \in \mathbb{R}^{m_g} : \frac{\|(\mathbf{g} * \mathbf{h})_{(m:m+m_1-1)}\|^2}{\|(\mathbf{g} * \mathbf{h})\|^2} \geq \alpha \right\} \quad (3)$$

the impulse response of the SCISF must belong to S^m for some m .

Let $\omega_1, \dots, \omega_N$ be the location of the sub-carriers in the frequency domain. The spectral constraint may be applied by selecting $\mathbf{g} \in S^m$, so that the cost function

$$J(\mathbf{g}) \triangleq \sum_{i=1}^N (|G(\omega_i)| - G_d(\omega_i))^2 \quad (4)$$

is minimized for the desired target spectral response $G_d(\omega)$.

Typically, this optimization would have to be repeated for all possible m to select the filter impulse response, \mathbf{g} , that achieves the lowest possible value of J . However, as discussed below, optimization may be limited to a few well chosen values of m .

The determination of the filter impulse response may be done in two stages. First, the desired magnitude frequency response $G_d(\omega)$ of the impulse shortening filter \mathbf{g} over the bins used by the DMT system is obtained. Second, J is optimized over S^m for a specific value of m .

To determine the target spectral response $G_d(\omega)$, it is necessary to use expressions for the signal to noise ratios observed in the various frequency bins at the output of the DFT in the receiver. The DMT system has M tones (subcarriers), N of which (those from M_1 through M_2) are used, and the communication channel has an analog frequency response $H_c(f)$. The analog noise power spectral density observed at the input of the receiver A/D is $S_n(f)$. Prior to conversion in the A/D converter, the received analog signal may be filtered by an anti-aliasing filter with transfer function $H_a(f)$. The EDIR in the absence of the impulse shortening filter is $h(n)$. After the A/D converter, the signal is fed into the impulse shortening filter with an impulse response of $g(n)$. The impulse shortening filter ensures that $(h(n)*g(n))$ is shorter in length than the cyclic prefix. $G(\omega)$ is the discrete time Fourier transform of $g(n)$. Under these

conditions, the expected signal energy $\mu(k)$ observed in bin k at the output of the length- $2M$ receiver DFT is given by:

$$\mu(k) = C_1 D_k \left| H\left(\frac{\pi k}{M}\right) \right|^2 \left| G\left(\frac{\pi k}{M}\right) \right|^2 ; H(\omega) = H_c\left(\frac{\omega}{2\pi T}\right) H_a\left(\frac{\omega}{2\pi T}\right) \quad (5)$$

where C_1 is a constant, $1/T$ the sampling frequency and D_k the transmitted power in bin k . The noise power $\eta(k)$ in bin k is:

$$\eta(k) = C_2 \left[S_\eta\left(\frac{\omega}{2\pi T}\right) \left| G(\omega) \right|^2 \left| H_a\left(\frac{\omega}{2\pi T}\right) \right|^2 \right] * \left[\frac{\sin^2(M\omega)}{\sin^2\left(\frac{\omega}{2}\right)} \right] \Big|_{\omega=\frac{\pi k}{M}} \quad (6)$$

where C_2 is another constant and $*$ denotes the convolution. Assuming that the noise in the bands corresponding to unused tones is sufficiently attenuated by the anti-alias filter, $\eta(k)$ is approximately equal to:

$$\eta(k) \approx C_3 \sum_{l=M_1}^{M_2} S_\eta\left(\frac{l}{2MT}\right) \left| G\left(\frac{\pi l}{M}\right) \right|^2 \left| H_a\left(\frac{l}{2MT}\right) \right|^2 (\alpha(k-l) + \alpha(2M-k-l)), \quad (7)$$

where $\alpha(n)$ is defined as:

$$\alpha(n) = \int_{-\frac{\pi}{2M}}^{\frac{\pi}{2M}} \left[\frac{\sin^2\left(M\left(\frac{\pi}{M} - \lambda\right)\right)}{\sin^2\left(\frac{1}{2}\left(\frac{\pi}{M} - \lambda\right)\right)} \right] d\lambda, \quad (8)$$

$M_1 \dots M_2$ are the used tones and C_3 another constant. Defining \mathbf{x} to be the vector of frequency magnitudes to be solved for as:

$$\mathbf{x} \triangleq \begin{bmatrix} \left| G\left(\frac{\pi M_1}{M}\right) \right|^2 \\ \vdots \\ \left| G\left(\frac{\pi M_2}{M}\right) \right|^2 \end{bmatrix}, \quad (9)$$

the SNR in bin k can be seen to be of the form $\frac{b_k x_k}{a_k x}$, in which b_k are scalars and a_k are vectors.

To determine the target spectral response, \mathbf{x} is chosen to maximize the bit throughput. Approximating the capacity of bin k

by $\sqrt{\text{SNR}(k)}$, the optimal spectral profile is obtained by minimizing the cost function F , where:

$$F(\mathbf{x}) = - \sum_{k=M_1}^{M_2} \sqrt{\frac{b_k x_k}{a_k^T \mathbf{x}}} . \quad (10)$$

This minimization is performed over:

$$X = \{ \mathbf{x} \in \mathbb{R}^N : \|\mathbf{x}\| = 1, x_i \geq 0, 1 \leq i \leq N \} \quad (11)$$

and can be accomplished by any one of a variety of constrained optimization strategies, as discussed in Bertsekas, Nonlinear Programming, Athena Scientific, Belmont, MA, 1995. A median filter may be applied to the output of the optimization algorithm to smooth the resulting target spectral response.

Let $\mathbf{A} \triangleq \mathbf{H}_{(0:m_h+m_g-2)}^{m_g}$ and $\mathbf{B} \triangleq \mathbf{H}_{(m:m+m_l-1)}^{m_l}$. The energy constraint in equation (3) can be written as:

$$S^m = \left\{ \mathbf{g} \in \mathbb{R}^{m_g} : \frac{\mathbf{g}^T \mathbf{B}^T \mathbf{B} \mathbf{g}}{\mathbf{g}^T \mathbf{A}^T \mathbf{A} \mathbf{g}} \geq \alpha \right\} . \quad (12)$$

Matrix \mathbf{A} has full column rank, since it corresponds to a full convolution, so $\mathbf{R}_A \triangleq \mathbf{A}^T \mathbf{A}$ is invertible. Let $\mathbf{R}_B \triangleq \mathbf{B}^T \mathbf{B}$.

Defining $\mathbf{q} = \mathbf{R}_A^{0.5} \mathbf{h}$, where $\mathbf{R}_A^{0.5}$ is the square root of the positive definite matrix \mathbf{R}_A , the energy constraint set can be written in terms of \mathbf{q} as:

$$\mathcal{Q} \triangleq \left\{ \mathbf{q} \in \mathbb{R}^{m_g} : \frac{\mathbf{q}^T \mathbf{R}_A^{-0.5} \mathbf{R}_B \mathbf{R}_A^{-0.5} \mathbf{q}}{\mathbf{q}^T \mathbf{q}} \geq \alpha \right\} . \quad (13)$$

The next step is to reduce the dimensionality, i.e., the number of variables to search over in the optimization process. For example, in a video digital subscriber line (VDSL) application, an impulse shortening filter having a few hundred taps may be required. Searching over a variable space so large is difficult and impractical. Instead, the optimization is

performed by searching over a cleverly chosen lower-dimensional subset of the variable space. This simplification in the optimization process may be done without significant reduction in the performance of the communication system.

5 The reduction in dimensionality is accomplished by a transformation of variables. Let $C \triangleq R_A^{-0.5} R_B R_A^{-0.5}$. Let $S = U \Sigma U^T$ be the singular-value decomposition of C , where Σ is a diagonal matrix whose (i, i) th element is σ_i and where the σ_i are arranged in descending order. If $\sigma_1 < \alpha$, there is no feasible solution
10 corresponding to delay m . If $\sigma_1 \geq \alpha$, let

$$U = [U_1 U_2] \quad \Sigma = \begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix}, \quad (14)$$

where U_1 has size (m_g, m_d) and Σ_1 has size (m_d, m_d) for some m_d . These equations define the dimension of the range space of the matrix U_1 , over which the search is confined. The dimension m_d
15 may be chosen either to include all σ_i greater than some threshold $\beta, \beta < \alpha$, or m_d may be a fixed number. Simulations indicate that m_d can be less than m_g by more than an order of magnitude without significantly affecting communication system performance.

20 The dimensionality reduction is achieved by a further transformation of variables: $q = U_1 v$. The energy constraint set now becomes:

$$V \triangleq \left\{ v \in \mathbb{R}^{m_d} : \frac{v^T \Sigma_1 v}{v^T v} \geq \alpha \right\} = \left\{ v \in \mathbb{R}^{m_d} : v \neq 0, \alpha \|v\|^2 - v^T \Sigma_1 v \leq 0 \right\} \quad (15)$$

using the identity $U_1^T C U_1 = \Sigma_1$.

25 Next, the cost function is expressed in terms of v . For a particular v , the corresponding g is given by $g = R_A^{-0.5} U_1 v$, which leads to:

$$\begin{bmatrix} G(\omega_1) \\ \vdots \\ G(\omega_N) \end{bmatrix} = \mathbf{F}\mathbf{g}; \quad \mathbf{F} = \begin{bmatrix} 1 & e^{-j\omega_1} & e^{-j2\omega_1} & \dots & e^{-j(m_g-1)\omega_1} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & e^{-j\omega_N} & e^{-j2\omega_N} & \dots & e^{-j(m_g-1)\omega_N} \end{bmatrix} \quad (16)$$

Let $\mathbf{D} \triangleq \mathbf{FAR}_A^{-0.5} \mathbf{U}_1$. Let $\mathbf{D}^R \triangleq \text{real}(\mathbf{D})$ and $\mathbf{D}^I \triangleq \text{imag}(\mathbf{D})$. \mathbf{D}^R and \mathbf{D}^I are real matrices of size (N, m_d) . Let $\mathbf{d}_{R,n}^T, 1 \leq n \leq N$ and $\mathbf{d}_{I,n}^T, 1 \leq n \leq N$ be the rows of \mathbf{D}^R and \mathbf{D}^I , respectively. Then:

$$\begin{aligned} |G(\omega_n)| &= \sqrt{(\mathbf{d}_{R,n}^T \mathbf{v})^2 + (\mathbf{d}_{I,n}^T \mathbf{v})^2} \\ &= \sqrt{\mathbf{v}^T (\mathbf{d}_{R,n} \mathbf{d}_{R,n}^T + \mathbf{d}_{I,n} \mathbf{d}_{I,n}^T) \mathbf{v}} = \sqrt{\mathbf{v}^T \mathbf{T}_n \mathbf{v}} \end{aligned} \quad (17)$$

where $\mathbf{T}_n \triangleq \mathbf{d}_{R,n} \mathbf{d}_{R,n}^T + \mathbf{d}_{I,n} \mathbf{d}_{I,n}^T$. These definitions result in:

$$J_o(\mathbf{v}) \triangleq j(g(\mathbf{v})) = \sum_{n=1}^N \left(\sqrt{\mathbf{v}^T \mathbf{T}_n \mathbf{v}} - G_d(\omega_n) \right)^2 \quad (18)$$

The gradient and Hessian of J_o are:

$$\nabla_{J_o}(\mathbf{v}) = \sum_{n=1}^N 2 \left(1 - \frac{G_d(\omega_n)}{\sqrt{\mathbf{v}^T \mathbf{T}_n \mathbf{v}}} \right) \mathbf{T}_n \mathbf{v} \quad (19)$$

$$\mathbf{H}_{J_o}(\mathbf{v}) = \sum_{n=1}^N 2 \left[\left(1 - \frac{G_d(\omega_n)}{\sqrt{\mathbf{v}^T \mathbf{T}_n \mathbf{v}}} \right) \mathbf{T}_n + \frac{G_d(\omega_n)}{(\sqrt{\mathbf{v}^T \mathbf{T}_n \mathbf{v}})^3} \mathbf{T}_n \mathbf{v} \mathbf{v}^T \mathbf{T}_n \right] \quad (20)$$

The projection $P_v(\mathbf{y})$ of any $\mathbf{y} \in \mathbb{R}^{m_d}$ on to V is defined to be

$$P_v(\mathbf{y}) \triangleq \arg \min_{\mathbf{v} \in V} \|\mathbf{y} - \mathbf{v}\|^2 \quad (21)$$

To develop an algorithm to optimize the cost function J_o over V , an expression must be derived for $P_v(\mathbf{y})$. There is no closed-form expression for this projection operator, however there is a very efficient algorithm to compute it. If $\mathbf{y} \in V$, $P_v(\mathbf{y}) = \mathbf{y}$. If not, the projection onto V is the same as the projection on to its boundary, defined by:

$$\text{boundary}(V) \triangleq \left\{ \mathbf{v} \in \mathbb{R}^{m_d} : \mathbf{v} \neq 0, \alpha \|\mathbf{v}\|^2 - \mathbf{v}^T \Sigma_1 \mathbf{v} = 0 \right\} \quad (22)$$

The latter may be accomplished using LaGrange multipliers as follows. Defining a modified cost function S as:

$$S(\mathbf{v}) = \frac{1}{2} \|\mathbf{y} - \mathbf{v}\|^2 + \mu (\alpha \|\mathbf{v}\|^2 - \mathbf{v}^T \Sigma_1 \mathbf{v}) \quad (22)$$

The equation $\nabla S(\mathbf{v}) = 0$ must be solved on the boundary of V . This
5 reduces to the following simultaneous equations:

$$[(1 + \mu\alpha)\mathbf{I} - \mu\Sigma_1]\mathbf{v} = 0 \quad (23)$$

$$\alpha \|\mathbf{v}\|^2 - \mathbf{v}^T \Sigma_1 \mathbf{v} = 0 \quad (24)$$

Solving for \mathbf{v} in terms of μ from (23) and substituting in (24), the following equation is obtained for μ :

$$10 \quad \mathbf{y}^T \text{diag} \left(\frac{\alpha - \sigma_i}{[1 + \mu(\alpha - \sigma_i)]^2} \right) \mathbf{y} = 0 \quad (25)$$

Equation (25) may be written as a polynomial equation of order $2m_d$ in μ . The polynomial equation must be solved for the real and positive roots, which may be done using one of the variety of efficient root-finding algorithms in the literature. One of the
15 real, positive roots must be chosen by enumeration so that, when substituted into the expression for \mathbf{v} , it leads to the smallest $\|\mathbf{y} - \mathbf{v}\|^2$.

Since the gradient and Hessian of the cost function J_0 and the projection operator onto V are available, the optimization
20 can be done very efficiently. For example, the penalty method may be used, as described in Bertsekas, Nonlinear Programming, Athena Scientific, Belmont, MA, 1995, or an iterative strategy consisting of gradient descent followed by projection onto V . Both techniques have been tested and perform well.

25 Other embodiments are within the scope of the following claims.

What is claimed is: